

Using Market Segmentation to Obtain Plant-Specific Instruments: A Practical Application

Chad Syverson
University of Maryland

November 1999

Abstract

A lack of plant-specific instruments has hampered some empirical studies using economic microdata. I propose that market segmentation can be exploited to identify appropriate instruments containing across-plant variation. I outline an intuitive principle for identifying such instruments, and I explore their effectiveness with a prototype study: production function/returns to scale estimation. The instruments, which can be plausibly used in a large variety of empirical work, are also shown to be surprisingly flexible in applicability across industries. These factors imply that instrumental variables identified through market segmentation offer many opportunities for expanded microdata research activity.

This is preliminary work; comments are welcome. Please contact author before citing.

The research in this paper was conducted while the author was a research associate at the Center for Economic Studies, U.S. Bureau of the Census. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census or the Center for Economic Studies.

Using Market Segmentation to Obtain Plant-Specific Instruments: A Practical Application

I. The Market Segmentation Principle of Instrument Identification

The recent proliferation of available economic microdata has greatly benefitted nearly every field in economics. Questions once left only to theory are now quantitatively testable. Researchers are gaining new insights into long-standing puzzles. The aggregate dynamics of many economic variables have been decomposed into their fundamental components. Establishment-level data has doubtlessly offered economists excellent new opportunities to learn. Unfortunately, the increased availability of such data has not been accompanied by a similar increase of suitable plant-specific instrument series. Instruments are crucial to many empirical studies because they allow econometricians to obtain consistent estimators in the presence of simultaneity/endogeneity problems, and economic data is rife with endogenously determined variables. Simply measuring economic data at a finer level does not rid it of endogeneity, because individual economic actors still base many actions on influences that are either unknown or unmeasured. Furthermore, it is often impossible to empirically model these effects. The use of instrumental variables (IV) estimation is warranted in such cases. Because of the lack of plant-specific instruments, however, microdata researchers encountering endogeneity to this point have had to compromise by employing techniques that, while perhaps practical, are less theoretically appropriate than instrumental variables estimation.

I contend that market segmentation—geographic segmentation here specifically—can be exploited to identify establishment-level instrument series. The instruments which I propose have a wide array of potential empirical applications that span several economic fields. Examples include using plant data to estimate labor supply curves, output supply curves, factor substitution elasticities, returns to scale, and total factor productivity. Besides flexibility in the nature of empirical questions which the instruments it identifies can be used to address, the market segmentation principle has two additional important characteristics. First, its intuitiveness and generality allow instruments to be identified for plants in a host of industries. Second, because the specific form of market segmentation varies from industry to industry, the researcher can be reasonably assured that identified instruments are especially relevant to the

industry of interest. These combined factors imply that instruments identified through market segmentation can significantly expand the potential scope of empirical microdata studies. I test their suitability for one such application here: industry production function/returns to scale estimation.

Generally speaking, the suitability of a variable (or alternatively, a set of variables) to instrument for endogenous explanatory variables depends on two factors. The first requirement is that the instrument be correlated with the endogenous variables. Higher degrees of comovement improve instrument performance. The second is exogeneity; the instrument should have as little correlation with the residual term as possible, since the inconsistency and small-sample bias of an IV estimator go to zero as the instrument and the residual become orthogonal. If plant-level data are being used for estimation, there is a third criterion: a suitable instrument must exhibit some variation across plants to gain any additional identifying power from the plant data. Aggregate or even industry-wide series will not suffice. It is this third criterion which has caused many researchers who work with plant-level data to forsake the search for instruments as too difficult, if not hopeless. I believe, however, that careful theoretical consideration of market structures allows identification of plant-specific instruments that are widely applicable.

The key to identifying such instruments is recognizing how markets are segmented across the plants of interest. Market segmentation, for my purposes, refers to any way in which a seemingly industry- or economy-wide market is actually comprised of a collection of heterogeneous “local” market units. (“Local” does not necessarily imply that the market is segmented geographically, although that is certainly the case for many goods and industries.) That is, markets are segmented whenever there is some degree of plant-level separation in the industry's output or inputs markets. Recognizing such market heterogeneity allows identification of instrumental variables that will exhibit across-plant variation when measured along the same dimension as the segmentation is present.

To clarify the principle, consider a few examples. If plants in an industry face localized labor markets—geographically segmented labor markets, in other words—then local real wage measures are candidate instruments for plants' labor inputs, assuming wage-taking plants. If establishments face varying shadow prices of capital because of financial heterogeneity, proxies of credit-worthiness or liquidity from firm financial data could potentially instrument for plant-

level investment and/or capital inputs. In many intermediate goods industries, tight producer-supplier linkages exist; an upstream plant may ship a significant portion of its output to a very limited number of downstream buyers. When this is the case, activity data for the downstream plant(s) can be used to instrument for inputs in the upstream establishment. These latter two examples are cases of non-geographic market segmentation.

These illustrations, and indeed the application of the market segmentation principle in general, are subject to two important caveats. First, a researcher must be confident that any instruments still exhibit the properties of relevance and exogeneity at the establishment level; across-plant variation alone is not sufficient. Consider the above example of using local real wage measures as instruments for labor inputs. Such data would not be suitable instruments if plants or firms were believed to have labor-market power; the local wage would no longer be exogenous to the plant's input decisions. A second unfortunate reality is that data limitations often preclude practical implementation of what would be a superb instrument in theory. We may not have lists of buyer-supplier linkages in many industries, for instance.

Still, the exploitation of market heterogeneity to identify instruments holds considerable potential. The two caveats enumerated above can often be addressed effectively with some effort on the part of the empirical economist. For example, the requirement of plant-level instrument exogeneity requires prudent consideration of theory when selecting instruments.¹ If such issues are successfully accounted for, then we have powerful and flexible establishment-level input instruments at our disposal.

In the remainder of the paper, I will demonstrate the merits of the market segmentation method with a practical application. I choose production function/returns to scale estimation as the prototype study for my method for several reasons. First, these estimates are interesting to many economists in their own right; production functions are ubiquitous in economics, and a substantial literature exists which explores their implications regarding returns to scale.² A vast majority of the returns to scale literature uses either industry or aggregate data. Extension of this work to plant data, even if only for select industries, may offer additional insights. Second, the

¹ The relevance criterion can be more directly addressed because it is statistically testable.

² For examples of empirical work in returns to scale, see Hall (1990), Burnside et. al (1995), Basu (1996), Basu and Fernald (1997), and Basu and Kimball (1997).

current non-IV methods of production function estimation (some of which I describe below) suffer from shortcomings that good instruments can overcome. Third, I can test the breadth of practical applicability of the method through the choice of the industries in the study. That is, by including industries with varying degrees of market segmentation, I can explore the question of the extent of heterogeneity necessary to attain instrument relevance at the establishment level.

II. Production Function/Returns to Scale Estimation: Problems and Proposed Solutions

Production function estimation has a long history in empirical economics, beginning over 70 years ago with Cobb and Douglas (1928). Since then, the quantitative relationship between outputs and inputs has received substantial attention in the agricultural, industrial organization, and macroeconomic literatures. The availability of plant-level databases has sparked new interest in using microdata to obtain production function and returns to scale estimates. This has, in turn, generated a need for accurate estimation procedures when plant-specific data are used.

A naive researcher may estimate production functions simply by regressing outputs on some functional form of inputs using ordinary least squares (OLS) methods. However, as Marschak and Andrews (1944) first pointed out, simultaneity of productivity and inputs cause such methods to provide inconsistent estimators of production function parameters. Their argument has become more widely accepted in the years since, as evidenced by work such as Hoch (1962), Mundlak and Hoch (1965), and Mundlak (1996), among others. To see the nature of the OLS estimation difficulty, consider a simple Cobb-Douglas technology specification³:

$$Y_t = A_t L_t^{\gamma_l} K_t^{\gamma_k} E_t^{\gamma_e} M_t^{\gamma_m} \quad (1)$$

where Y_t is gross plant output at time t , and L_t , K_t , E_t , and M_t are the plant's labor, capital, energy, and material inputs, respectively. A_t is a coefficient that captures factor-neutral productivity in the plant's technology. Taking logarithms and using lower-case letters to denote logged values yields

$$y_t = \gamma_0 + \gamma_l l_t + \gamma_k k_t + \gamma_e e_t + \gamma_m m_t + \omega_t + \eta_t \quad (2)$$

³ I am suppressing plant subscripts here.

In (2), I have decomposed logged productivity (the natural logarithm of A_t) into three separate terms. The first, fl_0 , is the mean log productivity level over the observations. The other two terms, H_t and $\$t$, are mean-zero deviations from this constant, the difference between them being that H_t is observable to the plant while $\$t$ is not. The sum of fl_0 and H_t can be interpreted as productivity privately observed by the firm in period t . The value $\$t$ can be thought of either as the plant's measurement error of its own productivity or as an unforecastable innovation to productivity. Thus H_t is a state variable in the plant's decision functions, but $\$t$ is not. Of course, to the econometrician, productivity is not observable; while fl_0 can be estimated, H_t and $\$t$ must be incorporated into the error term of any estimable specification of (2).

The incorporation of unobserved productivity H_t into the error term leads to biases in OLS estimates. Because changes in productivity alter the marginal product of inputs, the demand function for adjustable inputs (such as labor, energy, and materials in the production function above) will include the contemporaneous realization of productivity H_t ; i.e., $l_t = l_t(4H_t)$, and likewise for e_t and m_t . Hence the error term will be correlated with labor, energy, and materials inputs, causing the estimates of fl_l , fl_e , and fl_m to be biased. Furthermore, if H_t is serially correlated, today's productivity conveys information about future productivity, and the demand function for quasi-fixed inputs such as capital also includes H_t because plants adjust their input stocks to anticipated changes in their marginal productivities. This similarly leads to biased estimates of fl_k .

Responses in the empirical literature to the problem of input endogeneity have varied. One strategy used with establishment-level panel data is estimation using establishment fixed effects.⁴ The intuition of this approach is that any permanent differences in average productivities across plants can be removed by allowing the estimated constant to vary between them. In effect, it makes the presumption that H_t is constant through time for any given establishment, although its level varies cross sectionally. As seen in (2), if H is indeed constant, there is no correlation between the error term and inputs (recall that the error term $\$t$ is the innovation to productivity not incorporated by the plant into its factor demand decisions). Unfortunately, while the fixed-effects strategy does eliminate the inputs-productivity correlation

⁴ For applications of this technique, see Hoch (1962), Griliches (1980), and Harrison (1994).

in the cross section, it requires the unlikely assumption that there is no intertemporal within-plant productivity movement, leaving the door open to bias through correlation across time periods.

In a recent paper, Olley and Pakes (1996) take a different angle toward eliminating simultaneity biases. They advocate including a proxy to control explicitly for productivity in production function specifications. Their method is a three-step algorithm that uses observed plant variables and an assortment of standard techniques to create productivity proxies. It has in a very short time since become a popular method for estimating production functions with plant level data because of its clever treatment of endogeneity and its relative ease of implementation.⁵ The thrust of their procedure is the inversion of the plant-level investment function to back out a productivity proxy polynomial that contains only plant observables. They demonstrate such maneuvering is mathematically consistent if plant investment is a monotonically increasing function of plant productivity, and if productivity is the only unobserved establishment-specific variable in the investment function. The latter is an especially strong assumption, and I argue in Syverson (1999) that it does not often hold in practice. I demonstrate that when other plant-specific state variables do affect investment, the Olley-Pakes (O-P) algorithm provides biased estimates of production function parameters.⁶

One particular case where the Olley-Pakes method is liable, and where the use of appropriate instrumental variables would improve the accuracy of estimation, is when output markets are geographically segmented; i.e., when establishments sell a majority of their output to buyers in their immediate vicinities. Local markets can yield considerable spatial demand variation. As such, plants in locally focused industries are likely to take their idiosyncratic demand state into account when hiring inputs; demand (or expected demand) is thus an additional plant-specific variable in the input demand functions of these plants. I contend such a case is present in the industries in this study.

Instrumental variables techniques are a preferred alternative when the O-P method is likely to yield biased estimators. In practice, however, obtaining good instruments for plant-level

⁵ For examples of its application, see Griliches and Mairesse (1995), Aw, Chen, and Roberts (1997), Pavcnik (1998), and Levinsohn and Petrin (1999).

⁶ Olley and Pakes also make a similar assumption about the character of a plant's produce/liquidate decision which I contend can also lead to biases. This point is tangential to the discussion here, however, so I will not address it further. An interested reader should see Syverson (1999).

production data can be a challenging task. (Indeed, the call for methods such as Olley and Pakes' algorithm grew out of a perceived lack of quality instruments.) In the specific case of production function/returns to scale estimation, good instruments should be correlated with plant-specific inputs (e.g., employment and capital stock) but uncorrelated with productivity movements. The market segmentation principle is well suited to find just such instruments; I apply the method below.

III. Production Function/Returns to Scale Estimation Using Local Downstream Activity Measures as Instruments

The specific manufacturers in my prototype study are those plants in SIC industries 2611 (Pulp Mills), 2621 (Paper Mills), 2631 (Paperboard Mills), 3271 (Concrete Block and Brick), 3272 (Concrete Products Except Block and Brick), 3273 (Ready-Mix Concrete), and 3531 (Construction Machinery and Equipment). These industries were selected in part because all have a significant amount of their output used by firms in particular downstream sectors, without any upstream industry's output singularly accounting for a large portion of downstream costs. In the case the concrete and construction equipment industries, the downstream output purchaser is the construction sector. It is the finance, insurance, and real estate (FIRE) sector for the pulp and paper industries. As Shea (1993) argues, these characteristics make measures of construction or FIRE activity suitable instruments for inputs in these lines of business, at least at the industry level. My technique extends these instruments to the plant level by matching local construction or FIRE data to upstream-industry plants in the same geographic market. However, while these industries all fit well into Shea's framework for instrument selection, they may vary in suitability for my method. This is because the identification principle becomes more powerful as the extent of segmentation increases; in this case, as plants sell larger shares of their output within a locality. This heterogeneity is present to different degrees in the industries here. Careful examination of the results obtained with my methodology across differentially segmented industries will shed light on the extent of its practical application.

I clarify the intuition behind selection of local construction and FIRE activity measures as input instruments for my industries with the following discussion. Its specifics focus on the suitability of local construction activity measures as instruments for its corresponding upstream

industries, but the logic extends analogously to the linking of the pulp and paper industries and local FIRE activity data.

Construction is relevant to the productive scale (i.e., the input levels) of concrete and construction equipment plants because large portions of the output of these industries are used in final construction output. So construction activity and inputs in these plants are very likely to move together.⁷ Furthermore, because construction projects generally require output from a wide array of industries, the percentage of total costs of final construction firms attributed to ready-mix concrete or equipment alone is likely to be relatively small.⁸ This small cost share makes it less likely that any productivity advances in the concrete or construction equipment industries—which lower the relative cost of concrete or equipment—will alter the amount of construction activity, because idiosyncratic price drops in a single intermediate input will not greatly lower the total costs faced by final construction firms. Therefore, productivity movements in these industries are nearly (if not entirely) uncorrelated with final construction activity, satisfying the exogeneity criterion.⁹

It is with regard to the requirement that the instrument should exhibit interplant variation where the suitability of local construction activity as an instrument may differ between plants in the concrete and construction equipment industries. Why is this so? Because the degree to which plants incorporate their local output demand into production input decisions is based on the extent of geographic market segmentation in their industry's output market. More establishment-specific identification is afforded by using local (at the county level, say) construction activity instruments for plant inputs in locally focused industries than in industries with plants selling output across a larger area. For example, the high weight-to-value ratio of concrete makes it reasonable to assume that concrete plants sell the vast majority of their output locally, so local construction activity measures (which presumably reflect local demand for concrete) should be suitable plant-specific instruments. On the other hand, construction

⁷ For example, 86.0%, 73.4%, 79.8%, and 57.4% of the respective 1977 outputs of SICs 3271, 3272, 3273, and 3531 were purchased by firms engaged in new construction activity. See Shea (1992).

⁸ Looking at 1977 again, concrete and construction equipment output accounted for 6.5% and 10.5% of new construction costs that year, respectively.

⁹ The idea of using input-output linkages to identify instruments was proposed by Shea (1993). His paper offers a more thorough discussion of how one can identify instruments at the industry level which are both relevant and exogenous using demand and cost shares.

equipment is much more readily transportable, so there is less connection between an equipment plant's input decisions and construction activity in its local area.

Table 1 presents shipping distance data from the 1977 Commodity Transportation Survey for my industries. Such data allows interindustry comparison of the extent of output market localization, which can be reasonably measured as average shipping radius. The first column shows the percentage (by weight) of industry final output shipped less than 100 miles from the point of manufacture during the survey year. The second column shows the corresponding number for output shipped from 100 to 199 miles. These values are indicative of the degree of geographic output market segmentation in the industry. Notice that the concrete industries are extremely locally focused, while the construction equipment industry ships a substantial proportion of its output far away from its place of production.

The table implies we can be reasonably confident that construction activity in, for instance, Lancaster County, Nebraska (containing the city of Lincoln) will influence the input choices of a concrete plant in Lincoln, but not one in, say, Pima County, Arizona (containing Tucson). Conversely, fluctuations in Pima County's construction business will not affect the Lincoln plant. Therefore Lancaster (Pima) County construction activity can be used to instrument for inputs in a Lincoln (Tucson) concrete plant with reasonable assuredness that the relevance and exogeneity criteria are being met for each plant. If construction activity measures are spatially disaggregated enough, local activity measures will capture substantial interplant variance in the instrument series.

The table also indicates that I include regionally or nationally focused industries in the study. This facilitates exploration of the useful limits of geographic market heterogeneity as an instrument identifier. Just as the identification principle becomes more practical as markets become more segmented, it becomes less powerful as heterogeneity disappears. We may not be able to confidently assume that Lancaster County construction activity is very relevant to the input decisions of a construction equipment manufacturer in Lincoln, for instance. A more appropriate instrument choice in this regard might be nationwide construction activity. However, with such a measure we would of course be left with no variation between plants in the instrument. Including industries with varying degrees of output localization (as indicated by their shipment patterns) allows me to compare the relevance of local downstream instruments under

different degrees of segmentation.

IV. Data

Local Construction and FIRE Activity Data

The key to practical implementation of the market segmentation principle is instrument data which can be pared along the same axis as the market heterogeneity. In the present case, that requires data on construction and FIRE activity that can be measured at a geographically disaggregate level. Such data does exist. I use local construction and FIRE sector instruments derived from the Census Bureau's public-use County Business Patterns (CBP) annual data over the 1977-1993 period. The CBP contains summary information on the scale of economic activity by major industry for every county in the United States. There are two measures of economic activity for each surveyed industry: the number of employees during the March 12th pay period and the annual industry payroll. Public-use Census data at such a fine geographic resolution is often full of censored data, but this is a relatively minor obstacle in the case of the FIRE sector (SICs 60-67), and even less so in the construction sector (SICs 15-17). This is because the sectors' omnipresence and abundance of small firms allow full disclosure of summary statistics in all but the smallest of counties. For those counties with exact employment and payroll data withheld for the sake of confidentiality (roughly 9.5% of the over 50,000 county-year observations for FIRE, and 1.5% for construction), a total employment range is reported. In those cases, I simply use the mean of the range as the imputed employment for the period. I impute payroll for these observations by multiplying imputed employment by the corresponding sector's average per-employee payroll for that year, which is computed using data from full-disclosure counties. The impact of using imputed numbers is likely to be even less than their proportion indicates, as the typically small nondisclosure counties are less likely to contain sample plants in one of my industries. Real payrolls are constructed for each observation by dividing the reported nominal annual payroll by the same year's CPI value.

I estimate returns to scale using both employment and real payroll as plant input instruments because each measure has its own strengths and weaknesses. Total employment may be a more direct measure of construction activity than real annual payroll, for instance. However, because CBP data include employment numbers during only one pay period of the

year, CBP employment numbers are subject to measurement error. This is especially true in the construction industries, which have large intrayear fluctuations in employment and are very sensitive to exogenous factors such as the weather. Consistent seasonality would not be a problem; the (probably lower off-season) mid-March employment values would simply be a constant proportion of full (summer) sector employment in the county, so relative cross-sectional and intertemporal variations would be preserved, just at a smaller absolute scale. However, the construction sector in particular is subject to the fickle nature of the weather and other factors which can vary greatly across space and time. This causes the accuracy of mid-March employment as a measure of activity for the year to change idiosyncratically across counties and years. Using real annual payroll as an alternative reduces some of this noise, as this data measures activity over an entire year rather than trying to obtain an annual value with a small sample. Payroll is not without its own problems, however; it confounds the spatial differences in wages and worker skills with a pure activity measure. I estimate industry production functions using both measures for now, keeping in mind these relative differences when interpreting the results.

I also take advantage of the geographic dimension of the CBP survey to examine how changing the level of geographic aggregation of the construction or FIRE activity data affects the relevance of the instruments among industry groups. I aggregate the instrument data at three geographic levels. The finest aggregation is at the county level—as the data are originally reported. In this case, downstream construction or FIRE activity in a given county instruments for inputs at any sample establishments in that county. County activity is an extremely local measure, however, even for plants in locally-focused industries. While many such plants do likely operate largely within one county, it is also highly probable that a significant fraction sell their output outside the boundaries of their county. This is especially true for larger establishments in multi-county metropolitan areas, and in the Northeast, where counties are simply smaller in area than their western counterparts. Multicounty activity measures may be more appropriate in such instances. Using broader geographic instrument aggregates would also allow gauging of their effectiveness across industries with significantly different levels of localization. I therefore also instrument using construction activity data aggregated at two broader levels. The first, and the smaller of the two geographically speaking, is at the

Component Economic Area (CEA) level. The Bureau of Economic Analysis (BEA) creates CEAs by collecting together counties considered to be substantially intertwined economically and to share common dynamics. All counties in the U.S. are placed in a CEA; there are roughly 380 CEAs in the nation. The third and highest geographic instrument aggregate I use is at the Economic Area (EA) level. The BEA combines CEAs which are considered themselves to be economically interconnected into 172 EAs. Construction and FIRE sector employment and real payrolls for these larger geographic divisions are simply the sum of the county-level values for all counties within the CEA or EA. I do lose some across-plant variation in the instrument set when I go to the CEA or EA level, of course. The loss in identifying power may be a necessary trade-off in order to gain relevance in those industries with plants that largely operate beyond their county's borders.

Plant Level Production Data

I take plant output and inputs data from the Census Bureau's Longitudinal Research Database (LRD). The LRD is a longitudinally linked database of the establishments polled in the Annual Survey of Manufactures and the Census of Manufactures. It contains a wealth of information on plant production activity. Importantly here, it also contains the state and county where the establishment is physically located, so it is possible to match each plant with local instrument values at all three geographic aggregation levels. While annual LRD data is available from 1972 to 1995, my sample period was limited to 1980-92 on the front end because of availability limitations in the annual CBP instrument data, which is only available for 1977 onward (I require three lags of instrument values for each input observation), and on the back end because of limitations in external data sets used to merge in capital depreciation ratios, investment deflators, and the like.

The estimated production function is expressed in terms of gross output. Depending on the estimation method, plant inputs enter either separately as explanatory variables, or as a cost-share-weighted composite (I discuss this further below). Yearly nominal gross output is the plant's reported total value of shipments plus an adjustment for changes in inventories of final goods over the year. Nominal output is converted to a real value by dividing by an output price deflator for the plant's corresponding four-digit industry, taken from the Bartelsman-Becker-

Gray/NBER Productivity Database.

Plant-level labor inputs are the sum of production worker hours (a reported value in the LRD) and an imputed value for nonproduction worker hours. Nonproduction worker hours are constructed using the method of Davis and Haltiwanger (1991), where the number of nonproduction workers at the plant (the difference between reported total employment and the number of production workers) is multiplied by the average annual hours worked by nonproduction employees within the corresponding two-digit industry and year. These latter values are based on Current Population Survey data.

Real investment for each plant is calculated simply by dividing reported equipment and structures investments (the LRD contains separate capital data for each of the two capital types) by the respective Bureau of Labor Statistics (BLS) two-digit type-specific capital deflator.

I use a combination of two methods to construct the capital stocks for each plant. When data for a given plant is available in consecutive years, capital stocks (again, computed separately for equipment and structures) are constructed using the perpetual inventory method. I depreciate the previous period's capital stocks using BEA type-specific three-digit depreciation rates, and then add real investment values to obtain the current period's capital stocks. For plant-year observations not preceded by an observation of the same plant in the previous year (this includes a plant's first observation), I compute capital as the establishment's reported book value capital stock multiplied by the ratio of book to real values for the entire corresponding three-digit industry in that year. The industry-level capital stocks are from published BEA data. The value of any reported machinery or building rentals is inflated to a capital stock by dividing by the BLS's rental cost of capital series for the respective capital type. Finally, the capital stock used in production function estimation is constructed by summing the equipment and structures stocks.

Real materials usage is plant materials costs divided by a corresponding four-digit materials deflator. Energy input is the sum of electricity and fuel expenditures, deflated using a four-digit energy cost index. Each of the industry-specific price deflators used in this process is taken from the Bartlesman-Becker-Gray Productivity Database.

The input cost shares used to construct the composite input are computed as follows. Establishment labor costs are the sum of total salaries, wages, and benefits paid to permanent workers plus any costs from hiring contract labor. I compute capital costs as the product of

establishment capital stocks and the BLS capital rental cost series. Energy costs are the sum of electricity and fuel purchases, and materials costs are a separately reported item in the LRD. I sum the four to obtain total costs, and calculate shares using this value.

V. Methodology

I obtain returns to scale estimates from two production function specifications. The first is a simple logarithmic Cobb-Douglas production form:

$$y_{it} = \gamma_0 + \delta_t + \gamma_l l_{it} + \gamma_k k_{it} + \gamma_e e_{it} + \gamma_m m_{it} + \omega_{it} \quad (3)$$

where y_{it} is (logged) gross output of establishment i , and l_{it} , k_{it} , e_{it} , and m_{it} are measures of plant labor, capital, energy, and material inputs, respectively. The first term is the productivity level common to all industry plants and time periods. The second term, δ_t , is a time-specific constant which captures any overall industry productivity movements. The final term is plant-specific productivity.

The second specification follows the returns to scale literature and uses a cost-share-weighted composite input on the right hand side of the specification. That is, the production function is expressed as

$$y_{it} = \gamma_0 + \delta_t + \gamma_x x_{it} + \omega_{it} \quad (4)$$

where

$$x_{it} = s_l l_{it} + s_k k_{it} + s_e e_{it} + s_m m_{it}$$

and s_j is the plant-level cost share of input j . Under the assumption of cost minimization, the estimate of β_x is the degree of returns to scale.

The reasons for using two production function specifications arise from limitations in the alternative estimation methods. While the latter specification does not impose a functional form on the production function, the Olley-Pakes algorithm cannot be applied whenever flexible inputs (such as labor and materials) cannot be separated from quasi-fixed inputs (capital). Thus it is necessary to impose some structure when using their algorithm. I choose the Cobb-Douglas specification for the sake of simplicity. On the other hand, while instrumental variables estimates with market-segmentation instruments can be theoretically applied toward estimation

of the Cobb-Douglas specification, there is a practical consideration hampering such efforts. As Shea (1997) demonstrates, instruments should not only be relevant to each of the individual endogenous explanatory variables, they should have *linearly independent* relevance. This implies in this case that downstream activity measures should have an influence on each input (labor, capital, energy, and materials) that is independent of their influence on the other inputs. While some independence may be gained through the ability of the lag/lead structure of the instrument set to capture differing dynamic impacts across input demand functions, the high degree of comovement in the response of the inputs to downstream demand may overpower any such effect. This difficulty was realized in practice; attempts to estimate the Cobb-Douglas specification using IV methods yielded unacceptably high standard errors.¹⁰ The necessity of linearly independent relevance is obviously not an issue when using a composite input, so I estimate (4) using instrumental variables. As mentioned above, the composite input specification offers the further advantage of generality.

Returns to scale are estimated using observations from establishments in the LRD over the sample period; input coefficients are constrained to be equal for all plant-year observations within each industry. Industry groups are the four-digit groupings enumerated above, with one exception. I follow the practice of the BLS, which groups the two concrete products industries (3271 and 3272) together in its own productivity studies because their differentiation is driven more by product than by technology. Hence I analyze six distinct industries: pulp mills, paper mills, paperboard mills, concrete products, ready-mix concrete, and construction equipment.

OLS estimates are obtained under both specifications for use as a benchmark. I simply regress logged gross output on an intercept, twelve year dummies for 1981-92, and the input term(s). The year dummies allow estimation of β_0 .

The O-P estimation method is extensively detailed in Olley and Pakes (1996). Here, I merely summarize the mechanics of their estimation procedure as followed in this study. The

¹⁰ This does not mean that instruments obtained via market segmentation can only be applied to certain functional forms. The current restriction results from the fact my instrument set here, downstream demand shifts, tends to move several inputs simultaneously. If I could obtain additional instruments which have influence on specific inputs, such as the local real wage or capital cost measures discussed above, I could add these to the instrument set and likely gain linearly independent influence across inputs, allowing separate technology parameter estimation by input. Time and data constraints leave me to demonstrate the merits of my method using only downstream measures in the instrument set for now. I leave expansion of the set to future work.

first step of the process obtains fl_l , fl_e , and fl_m by regressing gross output (again all values are in logs) on an intercept; labor, energy, and materials inputs; and a fully-interacted fourth-order polynomial in capital and investment. Year dummies are included here as well. The estimated value of the capital-investment polynomial is an estimate of idiosyncratic plant productivity and is saved for use in the third and final stage of the estimation.

The second stage yields probit estimates of plant survival probabilities which are combined with the first-stage productivity estimates to proxy for expected productivity in the third stage. Before this step can be taken, the problem of defining plant survival in the data must be tackled. Survival determination is complicated somewhat by the nature of the LRD. Because it is comprised of plant observations from the quinquennial Census of Manufactures (CM) and the Annual Survey of Manufactures (ASM), which is comprised of five-year random-sample panels which allow for entry and exit, plants in the LRD can reappear and disappear not necessarily because they halt operations, but only because they are excluded from a particular ASM panel. Further, since the probability of selection for the ASM panel is proportional to the size of the establishment, industries with large numbers of small plants (such as concrete businesses) have many plant time series with missing years. This makes survival determination more complex than merely seeing if there is an observation in the next year. However, the O-P algorithm requires lagged variables in its third stage, greatly simplifying the sorting of plants into survivors and nonsurvivors. Given that consecutive-year observations are necessary, any plants not in an ASM panel are immediately removed from consideration. Furthermore, because my sample period ends in 1992 (a CM year), I know all of the establishments operating in the nation that year. Given these two factors, survival determination becomes a relatively simple matter of checking whether an observation in a given year is the last for a particular ASM plant. If a plant's final observation is in 1980-81, 1984-87, or 1989-92 (initial or intra-survey years), it is surely an exit and can be counted as non-survival.¹¹ Any plant-year observation with data in the following year for the same establishment is a survivor case. The classification of observations for 1983 and 1988 (which are the final years of their respective ASM panels) is a bit trickier. A final plant observation in 1983 (1988) can indicate exit only *between* 1983 and 1987 (1988 and

¹¹ While my sample period is from 1980-92, I have plant data for 1993. Hence I am able to determine if an ASM plant in the sample in 1992 survives to the following year or not.

1992) because a plant may disappear following these years simply as a result of being dropped from an ASM panel, not necessarily due to a halt in operations. Determining which exit scenario is true requires a check for the plant in the next CM year, 1987 (1992). So for exits in 1983 or 1988, I assume the plant stops production that year. This of course implies underestimation of the survival rate for these two years. However, the number of such plants is relatively small and is unlikely to significantly affect the estimated returns to scale.¹² A probit model of the binary survival indicator variable run on an interacted fourth-order investment-capital stock polynomial offers estimates of plant survival probabilities.

The final stage of the Olley-Pakes algorithm estimates the capital parameter. It takes as its dependent variable gross output minus the sum of labor, energy, and materials inputs multiplied by their respective first-stage coefficient estimates. This value is regressed on a constant, year dummies, lagged capital stock, and a fourth-order interacted polynomial in the one-year lags of the plant's estimated survival probability and productivity. All variables except survival probabilities are in logs. The estimation is nonlinear because the capital coefficient is present in the productivity estimate, which is part of the fourth-order polynomial. The estimate for this coefficient in this final stage is the estimated capital stock technology parameter fl_k .

I obtain instrumental variables estimates with standard two-stage least squares techniques, using local construction or FIRE activity measures as instruments. Each input observation is instrumented for by the current value, three lags, and one lead of the construction activity measure. This lag/lead pattern was chosen based on two considerations. The first is my prior belief about the extent of management decision horizons, both forward- and backward-looking. The second consideration is Buse's (1992) demonstration that superfluous instruments in an instrument set lead to estimation biases. The resulting lag/lead structure is a reconciliation of these two factors. Separate estimations are run for each instrument set.

VI. Results

¹² Olley and Pakes (1996) found very little change in their technology parameter estimates even when they entirely excluded the estimated survival probabilities from their procedure. I am taking not nearly so bold a step with my assumption; I am only slightly underestimating the probabilities rather than excluding them altogether. Hence I expect the simplifying assumption to have no impact on my estimates.

Instrument Relevance Tests

Table 2 presents the results of the IV first-stage regressions for the various instrument sets; these measure the degree of instrument relevance to the endogenous explanatory variables. In each case, I regress plant composite input on a constant, year dummies, and the lead, current, and three lags of the corresponding downstream activity measure. There are six instrument sets for each industry: downstream sector employment aggregated at the county, CEA, and EA level, and real sector payroll at the same three levels of aggregation. The expectation is that as the measures become more aggregated, the instruments will tend to become less relevant for the most local industries—like the concrete products and ready-mix industries—while becoming more relevant to the inputs of establishments in the construction equipment and paper industries, which tend to sell output throughout a much larger physical area. It is more difficult to predict a priori the relative effectiveness of employment and real payroll as instruments.

The table reports three relevance measures for each instrument set. The first is the total R^2 of the instrument set, including both the downstream activity terms and year dummies. The second is the marginal R^2 of the downstream activity measures; that is, the increase in the R^2 of the instrument set when local construction or FIRE sector activity measures are added to the time dummies. The third is the F-statistic for the null hypothesis that the five (lead, current, and three lags) downstream activity coefficients are jointly zero.¹³ The first two measures offer some perspective on the economic explanatory ability of the instruments, while the third is a rigorous statistical test of relevance. Including year dummies in the instrument set removes the effect of aggregate changes on plant-level inputs within an industry. Thus I am isolating in the downstream activity terms the effect of geographically idiosyncratic construction or FIRE activity movements on local inputs. Any predictive ability of the downstream activity measures

¹³ Wald tests using heteroskedasticity-robust standard errors yield qualitatively unchanged results; I do not report them here. Further, while plant-level productivity is almost surely has some degree of time-persistence, I do not report autocorrelation-robust standard errors and Wald statistics for either the first-stage or production function estimates for two reasons. First, practical implementation of the Newey-West consistent covariance estimator requires either a balanced panel of constant-frequency data, or that the window be set to the minimum number of observations on an individual plant in the panel. Because my sample contains several plants with only one observation and others with observations only every five years, implementing the Newey-West procedure with a window greater than zero (the degenerate case) does not make mathematical sense. Second, the very facts that the average panel length is so short compared to the cross-sectional dimension of the sample, and that there are several time periods between many same-plant observations in which persistence can die out, mitigate the variance-covariance matrix estimation error caused by autocorrelation.

found after taking out aggregate time effects bolsters my contention that local construction or FIRE activity is a good plant-level instrument for plants in the sample industries, because I am relying solely on the explanatory power of local downstream activity movements to determine local input choices.

The most striking feature of the table is the broad statistical relevance of the local downstream activity measures across nearly every industry and geographic aggregation level. This is true not only for the very local concrete plants, as we might expect, but also for the more widely operating pulp and paper plants. Even for those plants in the construction equipment industry, the least local industry in the sample, construction activity is statistically relevant at the 1% level at every geographic aggregation (although only marginally so using county level instruments). The evidence on economic relevance is more difficult to interpret because it is a subjective matter rather than a statistically testable hypothesis, but I find these results largely encouraging as well. Considering the substantial heterogeneity present in plant level data and the large number of plant-year observations, I believe that marginal R^2 values of local downstream activity measures in the 0.015-0.10 range, as found in some of my sample industries, are acceptable as evidence of economic influence. Larger values found in SICs 2631 and 3273 are certainly indicative of such effect, and the total R^2 of the instrument sets are also quite acceptable. The small marginal explanatory power of county construction activity on construction equipment plants suggests caution in interpreting any results using this instrument set, but this caveat is an exception and not the rule among the industries and instrument sets.

These results also accord with intuition regarding the changing relevance of the downstream activity instruments at different aggregations. In more locally operating industries, the instruments' economic and statistical relevance tends to decrease as these measures become more geographically aggregated. This is the case in concrete products, ready-mix, and pulp plants. This is not surprising, because we would expect that narrower aggregates more closely measure the demand conditions facing influencing input decisions in plants that sell most of their output nearby. Broader aggregations are still relevant, though less so, because of operation across county lines or comovement in county activity and activity in the corresponding CEA or EA. Conversely, relevance largely increases with instrument aggregation in the more broadly operating paper, paperboard, and construction equipment industries. This indicates that in these

industries, establishments respond more to demand shifts across wider operating regions. This is also as expected. It should be noted that even as relevance falls with changes in aggregation, the local downstream activity instruments are still quite relevant across the board.

There appears to be no appreciable difference between construction or FIRE employment and real payroll to explain upstream input movements. The relevance statistics are remarkably consistent across activity measures.

Returns to Scale Estimates

I report the returns to scale estimates from the various estimation methods and specifications in Table 3. The estimates and standard errors for equations (3) and (4) are shown in each industry's second and third rows. Both the Olley-Pakes and IV estimates differ consistently from their (likely biased) OLS counterparts. The IV estimates differ from the O-P values by varying degrees from industry to industry. Unfortunately, it is not possible here to discern how much of this difference is due to the more generalized IV specification and how much can results from biases present in the O-P estimates due to the omission of additional plant-level state variables. This makes direct empirical comparisons between the two methods' accuracy difficult. Instead, I appeal to theory for help interpreting the results.¹⁴

The relevance results offer conclusive evidence that local downstream activity is correlated with input movements in my sample industries. This indicates that it is very likely that establishments are considering the state of local demand when making input decisions. As I have shown in earlier work, we can expect the O-P method to yield biased estimates if plant management takes any establishment-specific state variables other than productivity and capital into account when making input choices. This bias is likely to manifest itself more the greater the effect of the additional state variable(s) on inputs, so its expected influence on the O-P estimates becomes larger as plants within an industry operate in smaller geographic areas and local demand conditions have greater sway over production. Hence the O-P estimates for the concrete industries are particularly suspect, and those in pulp, paper, and construction equipment

¹⁴ Note that even if Olley-Pakes and IV returns to scale estimates happen to be the same, they still may have different implications for plant level productivity estimates (which are the estimated residuals from the production function), because the product of the composite input and its coefficient will likely differ from the sum of the individual inputs and their respective coefficients.

to a lesser degree (although relevance evidence still indicates the probability of biases in these industries as well).¹⁵ It is this very relevance of local downstream activity to plant-level input movements, a cause of biases in the O-P method, that makes construction and FIRE activity measures good instruments for the plants in the sample. This fact and the strong case for their exogeneity imply favorable small sample bias and consistency behaviors in my IV estimators. Because of this, I am willing to put more credence in the accuracy of the IV returns to scale estimates, given that certain variable utilization considerations (discussed below) are addressed.

The O-P estimates imply smaller returns to scale than the corresponding OLS estimates in five of six industries. Theory is ambiguous regarding the expected direction of this movement; the change in the input coefficients depends on the extent and direction of covariance between the inputs and productivity, as well as the additional influence of the omitted state variable bias. The small O-P sample size relative to the OLS and IV values stems the requirement that the O-P algorithm requires lagged values in its third stage. Hence any single-year observations (such as those in CM years for plants not included in the corresponding ASM panel) must be dropped. There is no simple procedure to calculate a standard error for the returns-to-scale parameter (the sum of the input coefficients) for the O-P estimates because there one cannot obtain the covariance of the capital coefficient estimate with those for the flexible inputs, as their estimation occurs in completely different equations. Hence no standard errors are shown for these estimates.

I have reported instrumental variables estimates for each industry and for all six instrument sets. I do not eliminate the less relevant instrument sets from estimation after pretesting; Hall, Rudebusch, and Wilcox (1996) demonstrate doing so can lead to ex-post inconsistent estimates. We can, however, use the relevance measures as an ex-post guide to interpreting the results.

The IV estimates are larger than their OLS counterparts in most cases. The notable exception to this is the estimates for the construction equipment industry, especially those using county level instrument data. Of course, this is the instrument set which had the least relevance

¹⁵ I am ignoring for now other possibly plant-specific state variables, such as wages and effective capital costs, which may also affect input decisions in establishments within any of the industries, thereby introducing additional biases into the Olley-Pakes estimates.

to plant-level inputs, so this may explain part of this result. The direction of the movement from the OLS to IV coefficients depends upon the covariance between productivity and the composite input. If there is a positive correlation, then the OLS estimator will be biased upward and consistent IV estimates should be smaller. A negative correlation will cause IV estimates to be larger than the OLS estimates. The sign of the productivity-input covariance depends on many factors. If the industry is perfectly competitive and technology is factor-neutral, then increases in productivity will increase the marginal product of all inputs and establishments will purchase more of them, knowing that any increase in output will be bought up along a perfectly elastic demand curve. However, if an industry is not perfectly competitive (so that the firm cannot freely increase output without changing its price) and technology's influence varies across inputs, plants may react to productivity increases by changing their mix of inputs. The resulting change in the IV composite input parameter estimate from the OLS values is ambiguous; it depends on factor intensities and substitutabilities. So differences between the OLS and IV estimates across industries depend both on industry technologies and output markets. Further investigation into the specific sources, while interesting, are beyond the scope of this paper.

The precision of the IV estimates are within quite reasonable bounds, and not unlike those obtained using other estimation methods with plant data. They have larger standard errors than the corresponding OLS estimates, but this is expected since instrumental variables trades efficiency for consistency.

While a multitude of factors affect the results, there remains very little evidence for increasing returns in any of the estimates. In four of the six industries, the IV estimates indicate near-constant or slightly decreasing returns to scale. There is some evidence of slightly increasing returns to scale in the concrete industries, especially in concrete products, where estimates are statistically and economically greater than one for most instrument sets. The lukewarm evidence for increasing returns differs from results obtained in early attempts to estimate returns to scale with aggregate or industry-level data (e.g., Hall [1990]), which found strong indications of wide increasing demands. This evidence was later considerably weakened by correcting for factor utilization, as done by Burnside et. al (1995), Basu (1996), and others. The first set of estimates reported in table 3 does not explicitly correct for changes in factor

utilization yet still indicates only mild increasing returns.¹⁶ My hours-based labor input does account to some degree for labor utilization changes, but there is no adjustment margin in the utilization intensity of my measured capital input. So what evidence I do find for increasing in some industries may only result from the correlation between unmeasured plant-level utilization (which will be incorporated into the error term) and local downstream activity, because it is likely that these two variables will be positively correlated.

I explore this possibility by incorporating variable utilization into my specification. Basu and Kimball (1997) assert that adding the product of the plant's labor cost share and logged hours per worker (computed as the reported production hours divided by the number of production workers) to (4) is a general control for changes in both labor and capital utilization. They also show how reweight the composite input to build in Basu's (1996) assumption that materials cannot be substituted for value added. I estimate this specification for my industries using the same instrument sets and present the results in the industries' fourth and fifth rows in table 3.

Adjusting for utilization has little discernable effect on the estimates other than to increase their standard errors. The estimates for the industry showing the largest returns to scale in the initial estimations, concrete products, do fall noticeably and in a statistically significant manner, as expected. However, the changes in most of the other estimates (either up or down) are not statistically significant given the precision of the utilization-corrected estimates. It appears that the already weak evidence for increasing returns found in the constant-utilization specification is further thinned by controlling for input use on the intensive margin.

My returns to scale results have interesting implications regarding previous estimates. Given that returns to scale appear roughly constant in my plant data even without explicit capital utilization controls, it seems possible that the boldly increasing returns found in uncorrected aggregate and industry data were largely a function of aggregation biases. Basu and Fernald (1997) explore such effects sourced in the move from industry to aggregate data. While not necessarily a result of the process they describe, my results bolster their contention that data aggregation hides reallocation effects which are important to understanding plant-level processes.

In sum, the relevance tests are quite encouraging with regard to the flexibility of

¹⁶ Of course, it is possible that this is just an artifact of the small size of my industry sample. I may have just selected industries that happen to not have a large amount of increasing returns in their technologies.

instrumental variables obtained by recognizing segmented markets. This widens the scope of potential applications of the market segmentation method. Further, this study shows that applying such instruments toward investigation of interesting empirical questions is entirely feasible and can offer what are for the most part very reasonable results.

VII. Conclusions

I have argued for a new emphasis on instrument identification when working with endogenously determined plant level data. To present, many studies facing such data have either ignored the problem or made elaborate attempts to circumvent it because of a perceived difficulty in finding suitable instruments. In this paper, I outline an intuitive strategy for identifying instruments which are likely to be exogenous and relevant across plants. I go further by offering specific examples and employing them in a prototype empirical study. I find that local downstream demand measures are relevant across plants within an industry, at least for the industries in this study. Surprisingly, this is the case even among plants in industries which at first glance seem to operate nationally. This result is quite encouraging for its implications about the number of industries to which such instruments are applicable.¹⁷

Further, I wish to stress that instruments identified by exploiting market heterogeneity have additional applications not only across the number of industries, but also across the types of empirical questions which they can be used to address. These go well beyond just the simple returns to scale estimation in this study. Such instruments lend themselves well to analyses of a number of economic behaviors that require measures of exogenous demand fluctuations.

Empirical economists often wrestle with endogeneity problems. Having strategies for dealing with these difficulties in microdata, such as the market segmentation principle of instrument identification, will hopefully encourage work that was formerly not possible and allow improvement of previous studies.

¹⁷ Local construction activity data alone affords a significant number of potential industries to study. Because construction is such a significant part of GDP, establishments in many industries produce a substantial percentage (if not a majority) of their output to satisfy end uses in construction, and few compose a significant portion of total end costs. Moreover, many of these intermediate industries operate largely in their own local area. Thus local construction activity data are suitable for many industries beyond just the industries in this study. Just a few examples in manufacturing include the wood kitchen cabinets (SIC 2434), brick and structural clay tile (SIC 3251), and the fabricated structural metal (SIC 3441) industries.

References

- Aw, Bee-Yan; Chen, Xiaomin and Roberts, Mark. "Firm-Level Evidence on Productivity Differentials, Turnover and Exports in Taiwanese Manufacturing." NBER Working Paper no. 6235, October 1997.
- Basu, Susanto. "Procyclical Productivity: Increasing Returns or Cyclical Utilization?" *Quarterly Journal of Economics*, 111(3), 1996, pp. 719-751.
- Basu, Susanto and Fernald, John. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*, 105(2), 1997, pp. 249-83.
- Basu, Susanto and Kimball, Miles S. "Cyclical Productivity with Unobserved Input Variation." NBER Working Paper no. 5915, February 1997.
- Burnside, Craig; Eichenbaum, Martin and Rebelo, Sergio. "Capital Utilization and Returns to Scale." *NBER Macroeconomics Annual*, 1995, pp. 67-123.
- Cobb, Charles W. and Douglas, Paul H. "A Theory of Production." *American Economic Review*, 18(1), 1928, Supplement, pp. 139-72.
- Griliches, Zvi. "Returns to Research and Development Expenditures in the Private Sector." in J. W. Kendrick and B. Vaccara, eds., *New Developments in Productivity Measurement*, vol. 44 of *Studies in Income and Wealth*. Chicago: Chicago University Press, 1980, pp. 419-454.
- Griliches, Zvi and Mairesse, Jacques. "Production Functions: The Search for Identification." NBER Working Paper no. 5067, March 1995.
- Hall, Robert E. "Invariance Properties of Solow's Productivity Residual." in Peter Diamond, ed., *Growth, Productivity, and Employment*. Cambridge, MA: MIT Press, 1990, pp. 71-112.
- Haltiwanger, John; Foster, Lucia and Krizan, C.J. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." NBER Working Paper no. 6803, September 1998.
- Harrison, Ann E. "Productivity, Imperfect Competition and Trade Reform." *Journal of International Economics*, 36(1), 1994, pp. 53-73.
- Hoch, Irving. "Estimation of Production Function Parameters Combining Time-Series and Cross-Section Data." *Econometrica*, 30(1), 1962, pp. 34-53.
- Levinsohn, James and Petrin, Amil. "When Industries Become More Productive, Do Firms? Investigating Productivity Dynamics." NBER Working Paper no. 6893, January 1999.
- Marschak, Jacob and Andrews, William H. "Random Simultaneous Equations and the Theory of

- Production.” *Econometrica*, 12(3/4), 1944, pp. 143-205.
- Mundlak, Yair. “Production Function Estimation: Reviving the Primal.” *Econometrica*, 64(2), 1996, pp. 431-38.
- Mundlak, Yair and Hoch, Irving. “Consequences of Alternative Specifications of Cobb-Douglas Production Functions.” *Econometrica*, 33(4), 1965, pp. 814-28.
- Olley, Steven G. and Pakes, Ariel. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica*, 64(4), 1996, pp.1263-97.
- Pakes, Ariel. “Dynamic Structural Models, Problems and Prospects: Mixed Continuous Discrete Controls and Market Interactions,” in *Advances in Econometrics*, C. Sims (ed.) Cambridge: Cambridge University Press, 1994, pp. 171-259.
- Pavcnik, Nina. “Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants.” Mimeo, Princeton University, October 1998.
- Shea, John. “The Input-Output Approach to Instrument Selection: Extended Table III.” Mimeo, University of Wisconsin, Summer 1992.
- Shea, John. “The Input-Output Approach to Instrument Selection.” *Journal of Business and Economic Statistics*, 11(2), 1993, pp. 145-55.
- Shea, John. “Instrument Relevance in Multivariate Linear Models: A Simple Measure.” *Review of Economics and Statistics*, 79(2), 1997, pp. 348-52.
- Syverson, Chad. “Production Function Estimation with Plant-Level Data: Olley and Pakes or Instrumental Variables?” Mimeo, University of Maryland, April 1999.
- U.S. Bureau of the Census. *1977 Census of Transportation, Commodity Transportation Survey*.

Table 1: Percentage (by Weight) of Total Output Shipped by Distance

Industry	SIC	0-99 miles	100-199 miles
Pulp Mills	2611	32.8	13.1
Paper Mills	2621	15.7	10.8
Paperboard Mills	2631	21.0	11.3
Concrete Block and Brick	3271	94.2	5.2
Concrete Products, Ex. Block and Brick	3272	98.3	1.2
Ready-Mix Concrete	3273	94.4	1.7
Construction Equipment	3531	14.0	8.7

Source: 1977 Commodity Transportation Survey, Bureau of the Census

Table 2: Instrument Relevance Tests

Industry	SIC	Statistic	Instrument Set					
			County/Emp	County/Pay	CEA/Emp	CEA/Pay	EA/Emp	EA/Pay
Pulp Mills	2611	<i>N</i>	392	392	422	422	422	422
		Total R ²	0.284	0.287	0.136	0.137	0.17	0.166
		Marginal R ²	0.192	0.195	0.04	0.041	0.074	0.07
		F	20.09*	20.5*	3.798*	3.836*	7.254*	6.838*
Paper Mills	2621	<i>N</i>	2887	2887	3110	3110	3110	3110
		Total R ²	0.114	0.117	0.116	0.123	0.151	0.151
		Marginal R ²	0.073	0.076	0.077	0.084	0.112	0.112
		F	47.28*	48.84*	54.20*	59.86*	82.05*	81.98*
Paperboard Mills	2631	<i>N</i>	2081	2081	2219	2219	2219	2219
		Total R ²	0.163	0.164	0.14	0.146	0.191	0.191
		Marginal R ²	0.137	0.138	0.116	0.122	0.167	0.167
		F	67.62*	68.02*	59.21*	62.88*	90.61*	90.31*
Concrete Products	3271&2	<i>N</i>	15,363	15,363	15,892	15,892	15,892	15,892
		Total R ²	0.213	0.212	0.2	0.199	0.201	0.199
		Marginal R ²	0.038	0.037	0.013	0.012	0.014	0.012
		F	148.4*	143.0*	50.25*	46.54*	57.06*	49.99*
Ready-Mix Concrete	3273	<i>N</i>	20,988	20,988	21,787	21,787	21,787	21,787
		Total R ²	0.306	0.299	0.171	0.168	0.172	0.168
		Marginal R ²	0.197	0.19	0.051	0.048	0.052	0.048
		F	1191*	1136*	267.3*	254*	276.5*	254.5*
Construction Equipment	3531	<i>N</i>	4573	4573	4768	4768	4768	4768
		Total R ²	0.277	0.277	0.288	0.286	0.296	0.293
		Marginal R ²	0.002	0.002	0.009	0.007	0.018	0.015
		F	3.6*	2.709	12.53*	10.12*	23.45*	19.00*

Notes: This table presents results from regressions of plant inputs on the downstream activity instruments for the industries in the study. Statistics presented include the sample size, "*N*"; the total R² of the instrument set (both downstream measures and year dummies), "Total R²"; the increase in the R² when the downstream activity terms are added to the set, "Marginal R²"; and the F-statistic for joint significance of the five downstream instrument terms, "F". Results are shown for instrument sets comprised of one of two activity measures (employment or payroll) geographically aggregated at one of three levels (county, CEA, or EA). See text for details.

* Denotes Significance at the 1% Level

Table 3: Returns to Scale Estimates

Industry	Utilization	Estimation Method								
		OLS Separate	Olley-Pakes	OLS Composite	IV County/Emp	IV County/Pay	IV CEA/Emp	IV CEA/Pay	IV EA/Emp	IV EA/Pay
Pulp Mills	<i>N</i>	422	369	422	392	392	422	422	422	422
	Constant	1.007 (*)	0.909 (*)	0.904 (0.003)	1.081 (0.016)	1.108 (0.017)	1.126 (0.029)	1.177 (0.032)	1.146 (0.028)	1.192 (0.031)
	Variable			0.929 (0.003)	0.979 (0.044)	0.907 (0.082)	1.006 (0.044)	0.965 (0.061)	1.006 (0.047)	0.942 (0.071)
Paper Mills	<i>N</i>	3110	2741	3110	2887	2887	3110	3110	3110	3110
	Constant	0.998 (*)	0.959 (*)	0.927 (0.002)	0.991 (0.005)	1.004 (0.005)	1.007 (0.010)	1.041 (0.011)	1.042 (0.010)	1.073 (0.011)
	Variable			0.943 (0.002)	1.006 (0.019)	0.977 (0.050)	1.000 (0.015)	1.028 (0.010)	1.028 (0.013)	1.050 (0.009)
Paperboard Mills	<i>N</i>	2219	1923	2219	2081	2081	2219	2219	2219	2219
	Constant	1.011 (*)	0.984 (*)	0.948 (0.004)	0.672 (0.093)	0.551 (0.133)	0.871 (0.038)	0.835 (0.043)	0.941 (0.027)	0.916 (0.030)
	Variable			0.970 (0.003)	0.785 (0.122)	0.699 (0.182)	0.866 (0.048)	0.857 (0.056)	0.950 (0.040)	0.951 (0.057)
Concrete Products	<i>N</i>	15892	3621	15892	15363	15363	15892	15892	15892	15892
	Constant	1.017 (*)	0.900 (*)	0.885 (0.018)	0.931 (0.039)	0.947 (0.039)	0.828 (0.084)	0.835 (0.083)	0.959 (0.062)	0.953 (0.064)
	Variable			0.991 (0.013)	1.03 (0.091)	0.964 (0.095)	0.97 (0.088)	0.904 (0.184)	1.019 (0.180)	1.073 (0.058)
Ready-Mix	<i>N</i>	21787	4471	21787	20988	20988	21787	21787	21787	21787
	Constant	1.017 (*)	1.076 (*)	0.917 (0.005)	0.905 (0.020)	0.899 (0.020)	0.950 (0.019)	0.947 (0.018)	0.943 (0.016)	0.943 (0.016)
	Variable			1.001 (0.005)	1.401 (0.410)	0.837 (0.110)	0.979 (0.040)	0.964 (0.037)	0.908 (0.137)	0.982 (0.037)
Construction Equipment	<i>N</i>	4768	2183	4768	4573	4573	4768	4768	4768	4768
	Constant	1.006 (*)	0.889 (*)	0.853 (0.005)	0.884 (0.013)	0.881 (0.013)	0.925 (0.014)	0.919 (0.014)	0.904 (0.012)	0.903 (0.012)
	Variable			0.941 (0.005)	0.847 (0.111)	0.963 (0.038)	1.089 (0.193)	1.061 (0.057)	0.978 (0.035)	0.998 (0.039)

Notes: This table presents returns to scale estimates and their standard deviations from the various estimation methods. These include OLS with separate inputs, Olley-Pakes, OLS with a composite input, and IV with several instrument sets. No standard deviation could be computed for the Olley-Pakes estimates for reasons discussed in the text, so no standard deviation is shown for its OLS counterpart specification. The IV instrument sets differ by downstream activity measure (employment or payroll) and level of geographic aggregation (county, CEA, EA). IV estimates are presented for two specifications: one assuming constant input utilization, and one controlling for variable utilization. See text for details.